

METHODS AND APPARATUS FOR PROTECTION
AGAINST NETWORK FAILURES

5 **Related Application**

Sub A1 7 The present application is related to U.S. Patent Application Attorney Docket No. R. Nagarajan 12, filed concurrently herewith in the name of inventor R. Nagarajan and entitled "Efficient Architectures for Protection Against Network Failures," which is incorporated by reference herein.

10

Field of the Invention

The present invention relates generally to techniques for protecting an optical network or other type of network from a failure in a link, span, node or other element of the network, and more particularly to techniques which permit traffic to be redirected through the network in the event of such a failure.

Background of the Invention

Communication network technology is advancing at a rapid rate. For example, all-optical networks using wavelength division multiplexing (WDM) are currently being deployed for a wide variety of communication applications. WDM techniques allow optical signals having different wavelengths to be multiplexed into a single optical fiber. Each of the wavelengths serves as an optical carrier and can be used independently of the other wavelengths, such that different wavelengths may use different modulation formats to carry different signal types. In one possible implementation, each wavelength may carry a modulation signal representing a synchronous optical network/synchronous digital hierarchy (SONET/SDH) client payload, where each client is a SONET-rate time division multiplexed (TDM) application and the common carried signals are in an OC-x format, where "OC" denotes optical carrier and x denotes the rate, e.g., an OC-3 format, an OC-48 format, an OC-192 format, etc.

Such optical networks generally include routing elements such as wavelength switching cross-connects, wavelength adapters, wavelength interchanging cross-connects, etc. A wavelength switching cross-connect serves to cross-connect incoming wavelengths on a given input fiber to

different output fibers, but does not provide any transformation in wavelength. When only this type of routing device is present in an optical network, the network typically routes a given end-to-end demand using a single wavelength. If a primary network path assigned to the given demand fails, the demand generally must be carried on a secondary or restoration path using exactly the same wavelength as the primary path. A wavelength adapter is a device which allows conversion of wavelength at the client-network interface. A wavelength interchanging cross-connect is used to cross-connect incoming wavelengths onto different output fibers while also providing transformation of wavelengths.

One type of approach to providing failure protection in an optical network formed of these and other elements is to provide complete redundancy, such that the network includes a dedicated backup or secondary connection for each primary connection of the network. When a link, span or node of the primary connection fails, traffic may then be switched onto the corresponding elements of the secondary connection.

FIG. 1 shows an example traffic demand between two nodes A and Z of a given network. In this example, the demand is two units of OC- x traffic, each unit corresponding to one of the dashed lines between nodes A and Z.

FIG. 2 illustrates a conventional network protection approach as applied to the two units of OC- x traffic in the FIG. 1 example. This approach is an example of the above-noted complete redundancy approach, and is known in the art as “1+1” protection or “bridge and select.” In this approach, the first and second units of OC- x traffic are routed in the manner indicated at 100 and 110, respectively. More particularly, an original signal corresponding to the first unit of OC- x traffic is routed on a first link 102 between the nodes A and Z, while a copy of this signal is routed on a second link 104 between the nodes A and Z. A particular one of the original signal or the copy is then selected at the destination node Z. The second unit of OC- x traffic is routed in a similar manner, with an original signal corresponding to the second unit of OC- x traffic routed on a link 112 between the nodes A and Z, while a copy of this signal is routed on a link 114 between the nodes A and Z. Again, a particular one of the original signal or the copy is then selected at the destination node Z.

Sub A2
 5 The "bridge and select" approach illustrated in FIG. 2 provides complete redundancy for the capacity required to route the two units of OC-x traffic between nodes A and Z. However, this approach suffers from a number of significant drawbacks. For example, although the approach provides link protection, i.e., protection against a failure in one of the primary links 102 or 112, its fails to provide span protection, where span protection refers generally to an ability to switch locally from a primary trunk to a backup trunk. In addition, since a full unit of the OC-x traffic is assigned to each link, this approach does not accommodate preemptible traffic, and fails to provide any opportunity for quality of service (QoS) enhancement.

10 More sophisticated approaches may involve the use of a path restoration algorithm to provide automatic restoration of network traffic in the event of a primary path failure, while sharing restoration capacities whenever possible, so as to reduce the total amount of required redundant capacity.

Examples of known path restoration algorithms are described in, e.g., U.S. Patent No. 6,021,113 issued February 1, 2000 in the name of inventors Bharat T. Doshi et al. and entitled "Distributed Precomputation of Network Signal Paths with Table-Based Link Capacity Control," J. Anderson, B.T. Doshi, S. Dravida and P. Harshavardhana, "Fast Restoration of ATM Networks," JSAC 1991; W.D. Grover, "The Self-Healing Network: A Fast Distributed Restoration Technique for Networks Using Digital Cross Connect Machines," IEEE Globecom 1987; U.S. Patent No. 4,956,835, issued to W.D. Grover on September 11, 1990; C.H. Yang et al., "FITNESS: Failure Immunization Technology for Network Service Survivability," IEEE Globecom 1988; C. Edward Chow, J. Bicknell, S. McCaughey and S. Syed, "A Fast Distributed Network Restoration Algorithm," IEEE Globecom '93, pp. 261-267, 1993; and S. Hasegawa, Y. Okanone, T. Egawa and H. Sakauchi, "Control Algorithms of SONET Integrated Self-Healing Networks," and U.S. Patent Nos. 5,435,003 and 5,537,532, both entitled "Restoration in Communications Networks" and issued to R.S.K. Chng, C.P. Botham and M.C. Sinclair.

25 These more sophisticated approaches are often computationally intensive, and are therefore not appropriate or desirable in many applications. What is needed is a simple approach which utilizes redundancy but also overcomes the above-noted problems associated with the conventional "bridge and select" approach.

Summary of the Invention

The present invention provides techniques for protecting against network failures in an optical network or other type of network.

In accordance with the invention, units of OC- x traffic or other type of traffic are routed between nodes in a network on corresponding sets of trunks, such that the traffic is balanced between disjoint paths, and a restoration process for the traffic is implemented using service layer or transport layer switching. The service layer switching may be, e.g., Internet protocol (IP) switching or other type of service layer packet-based switching.

In a first illustrative embodiment of the invention, first and second nodes are connected by first and second sets of trunks, with each of the trunks in a given set of trunks supporting a designated portion of a given one of the units of traffic. For example, the units of traffic may be routed such that a first half of a given one of the units of traffic is routed on a first one of the trunks in a given one of the sets of trunks, and a second half of the given unit is routed on a second one of the trunks in the given set of trunks.

In other illustrative embodiments of the invention, the first and second nodes are connected by first and second sets of trunks so as to form a four-trunk ring, with each of the first and second sets of trunks including a primary trunk and a backup trunk. A given one of the units of traffic is then routed on either an upper or lower portion of the ring. For example, the given unit of traffic may be split equally between the primary trunk and the backup trunk associated with the upper or lower portion of the ring, or routed entirely on the primary trunk associated with the upper or lower portion of the ring.

The four trunk ring may be in the form of an IP/optical hybrid ring, in which case the restoration process is implemented using service layer switching, or a SONET/optical ring, in which case the restoration process is implemented using transport layer switching.

Advantageously, the present invention in one or more of the above-noted embodiments is able to accommodate preemptible traffic, and can provide an opportunity for quality of service (QoS) enhancement. The network failure protection techniques of the invention thus make more efficient use of redundant capacity than the conventional techniques described above. The techniques of the present invention are well-suited for use in complex fiber-based optical networks

which include wavelength select devices, wavelength adapters, wavelength interchange devices and other types of optical routers, but are more generally applicable to any other type of network and any type of transport medium.

These and other features and advantages of the present invention will become more apparent from the accompanying drawings and the following detailed description.

Brief Description of the Drawings

FIG. 1 shows an example traffic demand between two nodes of a network.

FIG. 2 illustrates a conventional "bridge and select" approach to routing of the example traffic of FIG. 1.

FIGS. 3, 4 and 5 show the routing of the example traffic of FIG. 1 in accordance with respective first, second and third illustrative embodiments of the invention.

FIG. 6 shows a block diagram of a network node in accordance with the invention.

Detailed Description of the Invention

The invention will be illustrated herein in conjunction with the routing of exemplary OC-x traffic between a pair of nodes in a network. It should be understood, however, that the invention is not limited to use with any particular type of traffic demand, network node or network, but is instead more generally applicable to any network traffic routing situation in which it is desirable to provide improved protection against network failures. For example, the network failure protection techniques of the invention may be utilized not only in optical networks, but also in telephone, cable and other electrical networks. The term "network" as used herein is therefore intended to include, e.g., optical networks, electrical networks and hybrid optical-electrical networks. The term "service layer" is intended to include without limitation Internet protocol (IP), asynchronous transfer mode (ATM), frame relay, or other type of communication technique associated with packet-based switching. The term "transport layer" is intended to include without limitation a layer which is below the service layer and provides a transport mechanism for services implemented in the service layer. Examples of transport layer communication techniques include synchronous optical network/synchronous digital hierarchy (SONET/SDH), optical networking, PDH, etc.

FIG. 3 shows the routing of the FIG. 1 traffic in accordance with a first illustrative embodiment of the invention. This embodiment utilizes an approach referred to herein as a 2F IP ring, where "2F" denotes "two fiber" and IP refers to Internet protocol. In this approach, the first and second units of OC-*x* traffic shown in FIG. 1 are routed in the manner indicated at 120 and 130, respectively. More particularly, the first unit of OC-*x* traffic is split in half and each half is routed on one of two separate OC-*x* trunks 122 and 124 between the nodes A and Z. The second unit of OC-*x* traffic is routed in a similar manner, i.e., the second unit is split in half and each half is routed on one of two separate OC-*x* trunks 132 and 134 between the nodes A and Z. Each of the trunks 122, 124, 132 and 134 has a capacity of one full unit of OC-*x* traffic, and is therefore loaded to 50% of its capacity by the routing shown in FIG. 3.

In the 2F IP ring approach illustrated in FIG. 3, the traffic is split between two disjoint paths. Failure detection is preferably implemented at the physical layer of the well-known Open Systems Interconnection (OSI) model. The physical layer of the OSI model is also referred to herein as the transport layer. Restoration is preferably implemented at the network layer of the OSI model. The IP layer referred to herein is an example of the network layer of the OSI model. A failure may be detected at the transport layer and communicated to the IP layer so as to force a switch of the one-half unit of OC-*x* traffic from the affected trunk to the alternate trunk. It should be noted that, although it is believed that no standard signaling exists for this type of cross-layer communication, numerous suitable implementations for such switching will be readily apparent to those skilled in the art.

Advantageously, since each OC-*x* trunk 122, 124, 132 and 134 is loaded to only 50% of its capacity, the 2F IP ring approach allows the implementation of enhanced Quality of Service (QoS) capabilities, such as reduced delay or packet loss, while there are no maintenance or restoration activities taking place.

FIG. 4 shows the routing of the FIG. 1 traffic in accordance with a second illustrative embodiment of the invention. This embodiment utilizes an approach referred to herein as a 4F IP/optical hybrid ring, where "4F" denotes "four fiber" and IP again refers to Internet protocol. The 4F IP/optical hybrid ring is shown generally at 150 in FIG. 4, and includes nodes A and Z, primary OC-*x* trunks 152 and 156 illustrated as solid lines, and backup OC-*x* trunks 154 and 158 illustrated

as dashed lines. The downward arrow 160 denotes a span switching operation, in this case from primary trunk 152 to backup trunk 154.

In this approach, the first and second units of OC- x traffic shown in FIG. 1 are routed such that one unit of OC- x traffic is routed on the top half of the ring 150, while the other unit is routed on the bottom half of the ring 150.

There are two options for implementing this type of loading. In the first loading option, a given unit of OC- x traffic is split equally between the primary and backup trunks on a particular half of the ring 150. For example, the first unit of the OC- x traffic may be split equally between the primary and backup trunks 152 and 154, respectively, while the second unit of the OC- x traffic may be split equally between the primary and backup trunks 156 and 158, respectively. This loading option allows enhanced QoS to be provided in a manner similar to that previously described in conjunction with the 2F IP approach. Each of the trunks 152, 154, 156 and 158 has a capacity of one full unit of OC- x traffic, and is therefore loaded to 50% of its capacity by this first option.

In the second option, the entire unit of OC- x traffic for a given half of the ring 150 is routed on the primary OC- x trunk for that half. For example, the entire first unit of OC- x traffic may be routed on primary trunk 152 in the upper half of the ring 150, while the entire second unit of OC- x traffic is routed on the primary trunk 156 in the lower half of the ring 150. The primary trunks 152 and 156 are thus loaded to 100% of their capacity in this example. As a result, it is not possible to provide enhanced QoS.

The 4F IP/optical hybrid ring approach illustrated in FIG. 4 thus balances the traffic between two disjoint paths. As in the 2F IP ring approach of FIG. 3, failure detection is again preferably implemented at the transport layer, while restoration is implemented at the IP layer.

In the first loading option described above, i.e., when the load is split equally between primary and backup trunks, span switching may be performed at the IP layer with the aid of transport layer signaling. In the second loading option, i.e., when the load is placed entirely on the primary trunks, a transport layer span switch can be used and no IP layer action is required. Note that in both the first and second options, multiple span maintenance activities on disjoint spans can proceed in parallel.

As noted above, detection of a node or fiber span failure may again be implemented at the transport layer. More specifically, such a failure may be detected at the transport layer and communicated to the IP layer so as to control the switching of OC-*x* traffic from a primary trunk to a backup trunk on a given half of the ring, or from one half of the ring to the other half of the ring.

5 It should be noted that the 4F IP/optical hybrid ring approach provides additional reliability relative to a conventional 2F optical ring in the presence of multiple failures. For example, when a span maintenance is in effect on the primary trunk of one half of the ring and the backup fiber is cut, the traffic can be switched to the other half of the ring. In the 2F optical ring case, the traffic is lost, as both halves of the ring are unavailable.

10 FIG. 5 shows the routing of the FIG. 1 traffic in accordance with a third illustrative embodiment of the invention. This embodiment utilizes an approach referred to herein as a 2 node and 4 fiber SONET/optical ring. The 2 node and 4 fiber SONET/optical ring is shown generally at 170 in FIG. 5, and includes nodes A and Z, primary OC-*x* trunks 172 and 176 illustrated as solid lines, and backup OC-*x* trunks 174 and 178 illustrated as dashed lines. The downward arrow 180 denotes a span switching operation, in this case from primary trunk 172 to backup trunk 174. The ring 170 further includes routers 190 and 192, coupled to nodes A and Z, respectively. The nodes A and Z in this embodiment are assumed to be implemented as add-drop multiplexers (ADMs), although other devices could also be used.

20 As in the 4F IP/optical hybrid ring approach illustrated in FIG. 4, the first and second units of OC-*x* traffic shown in FIG. 1 are routed in this approach such that one unit of OC-*x* traffic is routed on the top half of the ring 170, while the other unit is routed on the bottom half of the ring 170.

25 Like the FIG. 4 approach, the 2 node and 4 fiber SONET/optical ring approach of FIG. 5 balances the traffic between two disjoint paths. However, in the FIG. 5 approach, both failure detection and restoration are preferably implemented at the transport layer. In addition, span switching and restoration in the FIG. 5 approach may be performed in accordance with well-known conventional SONET/optical ring techniques.

There is generally no possibility for enhanced QoS in the FIG. 5 approach. However, if protection access is available, it may be possible to use the protection channels for preemptible traffic.

Advantageously, since the FIG. 5 approach performs both failure detection and restoration at the transport layer, it is likely to be the fastest in terms of restoration of the illustrative approaches described herein.

FIG. 6 shows an exemplary network node 200 which may be used to implement the above-described routing functions. The node 200 may therefore correspond, e.g., to node A or node Z in the illustrative embodiments of FIGS. 3 through 5.

The network node 200 includes a controller 210, a switch fabric 212, a first line card 214 having a set of OC-x ports 215 associated therewith, and a second line card 216 having a set of OC-x ports 217 associated therewith. It should be understood that the node 200 has been simplified for purposes of illustration. For example, the node 200 in practice may include a substantially larger number of line cards and ports, as required for a given application.

The controller 210 includes a processor 220 and a memory 222. The processor 220 may be, e.g., a microprocessor, a microcontroller, an application-specific integrated circuit (ASIC) or other type of processing devices, as well as portions or combinations of such devices. The memory 222 may include an electronic random access memory (RAM), a read-only memory (ROM) or other type or memory device, as well as portions or combinations of such devices. The memory 222 may be used to store a demand database for storing demands for network capacity, and a set of routing tables which specify routing paths through a corresponding network for particular demands.

It should be noted that the node 200 may be an element of an optical network or other type of network which includes a very large number of nodes, and possibly a central controller. One or more of the nodes and the central controller may each represent a computer, processor-based switch or other type of processor-based device configured to provide the network protection techniques described herein. The invention is well-suited for use in large-scale regional, national and international networks which may include many subnetworks, each having hundreds of nodes, but can be used in any network application.

5